

Unified Framework and Algorithm for Quantized Compressed Sensing

Zai Yang*, Lihua Xie*, and Cishen Zhang†

Abstract

Compressed sensing (CS) studies the recovery of high dimensional signals from their low dimensional linear measurements under a sparsity prior. This paper is focused on the CS problem with quantized measurements. There have been research results dealing with different scenarios including a single/multiple bits per measurement, noiseless/noisy environment, and an unsaturated/saturated quantizer. While the existing methods are only for one or more specific cases, this paper presents a framework to unify all the above mentioned scenarios of the quantized CS problem. Under the unified framework, a variational Bayesian inference based algorithm is proposed which is applicable to the signal recovery of different application cases. Numerical simulations are carried out to illustrate the improved signal recovery accuracy of the unified algorithm in comparison with state-of-the-art methods for both multiple and single bit CS problems.

1 Introduction

The recently developed compressed sensing (CS) theory and methods [1, 2] can achieve acquisition of information contained within a huge volume of data using only a small amount of measurement samples. Different from the classical Shannon-Nyquist sampling theorem which requires that the sampling frequency be twice as many as the bandwidth of a signal in order to reconstruct its complete information, the CS theory assesses the success of signal recovery with the sparsity. A signal $\mathbf{x} \in \mathbb{R}^N$ of length N is called K -sparse in a basis $\Psi \in \mathbb{R}^{N \times N}$ if all but at most a number of $K \ll N$ entries of its coefficient $\boldsymbol{\theta} \in \mathbb{R}^N$ are zeros with $\mathbf{x} = \Psi\boldsymbol{\theta}$. Without loss of generality we assume that Ψ is an identity matrix, i.e., \mathbf{x} is sparse in the canonical basis, since for a general basis Ψ it can be absorbed into the following introduced sensing matrix \mathbf{A} . Rather than observing directly the original sparse signal \mathbf{x} , a number of M , $K < M \ll N$, linear measurements are acquired in CS as

$$\mathbf{y} = \mathbf{A}\mathbf{x} + \mathbf{n} \quad (1)$$

where $\mathbf{y} \in \mathbb{R}^M$ is the measurement vector, $\mathbf{A} \in \mathbb{R}^{M \times N}$ denotes the sensing/measurement matrix and $\mathbf{n} \in \mathbb{R}^M$ is the measurement noise vector. Though the recovery of \mathbf{x} from \mathbf{y} is generally ill-posed (less linear equations than the unknown variables), it is shown in [3] that a sparse signal \mathbf{x} can be stably recovered under mild conditions on \mathbf{A} in the sense that the recovery error grows linearly with the noise level. To do this, the basis pursuit denoising (BPDN) problem

$$\min \|\tilde{\mathbf{x}}\|_1, \text{ subject to } \|\mathbf{y} - \mathbf{A}\tilde{\mathbf{x}}\|_2 \leq \epsilon \quad (2)$$

is solved where $\epsilon \geq \|\mathbf{n}\|_2$ denotes the noise level. The recovery is exact in the noise free case. A similar result holds for compressible signals that are not exactly sparse.

Sparse Bayesian learning (SBL) [4, 5] was derived from the research area of machine learning and has become a popular method for sparse signal recovery in CS. In SBL, the sparse signal recovery problem is formulated from a

*Z. Yang and L. Xie are with EXQUISITUS, Centre for E-City, School of Electrical and Electronic Engineering, Nanyang Technological University, 639798, Singapore (e-mail: yang0248@e.ntu.edu.sg; elhxie@ntu.edu.sg).

†C. Zhang is with the Faculty of Engineering and Industrial Sciences, Swinburne University of Technology, Hawthorn VIC 3122, Australia (e-mail: cishenzhang@swin.edu.au).

Bayesian perspective while the sparsity information is exploited by assuming a sparse prior for the sparse signal of interest. As an example, a Laplace prior corresponds to the ℓ_1 norm widely used in optimization approaches [5, 6]. Since the exact Bayesian inference is typically intractable, approximation approaches to Bayesian inference have been adopted including evidence procedure [7], e.g., in [6], and variance message passing (VMP) [8], e.g., in [9]. One merit of Bayesian CS is the flexibility of modeling sparse signals that can not only promote the sparsity of its solution, e.g., in [6], but also exploit additionally known structure of the sparse signal, e.g., in [10]. Since the Bayesian inference is a probabilistic method and based on heuristics to some extent, one shortcoming of Bayesian approaches is that there have been fewer results on their signal recovery accuracy in comparison with deterministic approaches, e.g., BPDN.

The conventional CS framework is mainly focused on the sparse signal recovery from the real-valued measurement \mathbf{y} that has infinite bit precision. The required number of measurements M is mainly studied for guaranteed signal recovery accuracy [11–14]. Since quantization is necessary for practical considerations, e.g., data storage and transmission, we study the sparse signal recovery from quantized measurements in this paper. During the quantization process, each continuous-valued measurement is quantized into some value in a finite set. A new challenge is thus the existence of quantization errors. The noise free case with a uniform unsaturated quantizer is studied in [3, 15, 16]. A solver with quantization consistency is recommended in [15] that corresponds to replacing the ℓ_2 norm in BPDN by the ℓ_∞ norm. A BPDN solver is used in [3] to handle the quantization errors as additive white Gaussian noises because of the stable signal recovery performance of BPDN. A family of solvers, named as basis pursuit dequantizer of moment p (BPDQ $_p$), that includes BPDN and that in [15] as special cases is studied in [16] where the ℓ_2 norm in BPDN is replaced by an ℓ_p norm with $2 \leq p \leq +\infty$. By characterizing the quantization errors as independent random variables uniformly distributed in a common interval, it is shown in [16] that the optimal signal recovery accuracy is obtained at some finite $p \geq 2$. But unfortunately, the optimal p cannot be explicitly given in practice. Note that the common uniform distribution assumption is crucial to obtain the results in [16]. As a result, it is unclear whether the results in [16] can be extended to a general quantizer case where such an assumption fails. It is obvious that both BPDN and BPDQ $_p$ are inappropriate in the case of a saturated quantizer since data saturation may lead to large or even unbounded quantization errors that deteriorate their performance. To deal with the data saturation, Laska *et al.* [17] propose two modified versions of BPDN, including saturation rejection and consistency approaches, denoted by BPDN-SR and BPDN-SC respectively. The saturated measurements are simply rejected in BPDN-SR while they are incorporated in the signal recovery process in BPDN-SC. While quantization errors and measurement noises are coupled in most existing methods (some methods, e.g., BPDQ $_p$, consider only the noise free case to avoid such a problem), e.g., in [17], they are separately studied by Zymnis *et al.* [18] where the authors seek to find a signal estimate that maximizes the likelihood of the quantized measurements while ℓ_1 norm is exploited to promote the signal sparsity. The resulting algorithm is quoted as ℓ_1 -regularized maximum likelihood (L1RML) algorithm.

An extreme case of quantized CS is so-called 1-bit CS where each quantized measurement keeps only the sign information of the real-valued measurement and thus uses just one bit. The 1-bit CS framework is proposed in [19] and has attracted many research interests because it possesses many merits. For example, a 1-bit quantizer is a simple comparator that tests whether the measurement is above or below zero, leading to an easy implementation and a fast quantization process. A measurement noise can be neglected in 1-bit CS as long as it does not change the sign of the measurement. It is shown in [20] that to acquire just one bit for each measurement is optimal in the presence of heavy noises. The 1-bit case is quite different from the multi-bit case since all measurements are saturated in 1-bit CS and the signal scaling information is lost. A common approach to the signal scaling problem is to impose that the signal to be recovered has a fixed unit norm and then search for the signal on the unit hyper-sphere rather than in the whole space. Such a constraint is nonconvex and brings new challenges to algorithm design. Existing algorithms based on this constraint include renormalized fixed point iteration (RFPI) [19], matching sign pursuit (MSP) [21], restricted-step shrinkage (RSS) [22] and binary iterative hard thresholding (BIHT) [23]. Convex formulations of the 1-bit CS problem have been recently proposed by Plan and Vershynin [24, 25]. They show in [24] that a linear program can decode the noiseless case with guaranteed signal recovery accuracy under similar mild conditions as in

conventional CS. In [25] they introduce a seemingly unrelated convex program for the noisy case and show similar results. It is noted that both the BIHT and the convex problem in [25] that deal with the noisy 1-bit CS problem require the signal sparsity information (BIHT needs to know $\|\mathbf{x}\|_0$ and CVXP requires a proper upper bound for $\|\mathbf{x}\|_1$).

In this paper, we introduce a unified framework for quantized CS that works for both multi- and 1-bit CS cases. The new framework deals with quantization errors and measurement noises separately, allows data saturation in multi-bit CS, and does not need the signal sparsity information. Based on Bayesian inference, we show that many existing formulations for quantized CS are special cases of the new framework or related to it after minor modifications. Based on the new problem formulation, we propose an algorithm within the Bayesian CS framework where a three-layer hierarchical prior introduced in [9] is adopted as the sparse signal prior and Bayesian inference is carried out using VMP. The performance of the proposed algorithm is studied by extensive numerical simulations in various scenarios. It is shown that the new algorithm improves the signal recovery accuracy in comparison with state-of-the-art methods in both multi- and 1-bit CS.

Notations used in this paper are as follows. Bold-face letters are reserved for vectors and matrices. For ease of exposition, we do not distinguish a random variable from its numerical value. x_i is the i th entry of a vector \mathbf{x} . $\mathbf{x}_{\mathcal{I}}$ denotes a truncated vector of \mathbf{x} with entry indices in a set \mathcal{I} . \mathbf{A}_i is the i th column of a matrix \mathbf{A} . $\|\mathbf{x}\|_0$ counts the number of nonzero entries of a vector \mathbf{x} . $\|\mathbf{x}\|_p = (\sum_i |x_i|^p)^{1/p}$ denotes the ℓ_p norm of a vector \mathbf{x} with $1 \leq p \leq +\infty$. $\langle g(x) \rangle_{p(x)}$ denotes the expectation of a function $g(x)$ with respect to a random variable x whose probability density function is $p(x)$. \odot denotes Hardamard (elementwise) product. \succcurlyeq and \preccurlyeq denote \geq and \leq respectively with an elementwise operation.

The rest of the paper is organized as follows. Section 2 introduces the unified framework for quantized CS and studies its relations with existing formulations. Section 3 introduces the proposed Q-VMP algorithm. Section 4 presents numerical simulations to illustrate the improved signal recovery accuracy of the proposed Q-VMP algorithm in comparison with existing ones. Section 5 concludes the paper and discusses some future works.

2 A Unified Framework for Quantized CS

2.1 A Unified Observation Model

In quantized CS, the observed samples are noisy linear measurements of the original signal after quantization:

$$\mathbf{z} = \mathcal{Q}(\mathbf{y}), \quad \mathbf{y} = \mathbf{A}\mathbf{x} + \mathbf{n} \quad (3)$$

where \mathbf{x} is the signal of interest, \mathbf{A} is the sensing matrix, \mathbf{n} is the measurement noise vector, \mathbf{y} is the pre-quantized noisy measurement vector, \mathcal{Q} denotes a quantizer and \mathbf{z} is the observation. A (regular) quantizer $\mathcal{Q}(v)$ for a scalar $v \in \mathbb{R}$ is defined as

$$\mathcal{Q}(v) = \begin{cases} v_0, & \text{if } v \in (u_0, u_1), \\ v_1, & \text{if } v \in [u_1, u_2), \\ \dots, & \dots, \\ v_{L-1}, & \text{if } v \in [u_{L-1}, u_L), \end{cases} \quad (4)$$

where L denotes the quantization level and typically satisfies $L = 2^B$ with B denoting the bit depth (bits per quantized measurement), $u_0 < u_1 < \dots < u_L$, and $v_i \in [u_i, u_{i+1})$ for $i = 0, \dots, L-1$. The quantizer $\mathcal{Q}(v)$ is called unsaturated if (u_0, u_L) is a finite interval, or saturated otherwise. For a vector \mathbf{v} , $\mathcal{Q}(\mathbf{v})$ operates elementwise. Multi-bit CS refers to the case $B \geq 2$ while 1-bit CS corresponds to $B = 1$.

2.1.1 Multi-bit CS

We consider first a multi-bit quantizer where $B \geq 2$. Denote \mathcal{D}_y the domain of \mathbf{y} . Then we have

$$\mathcal{D}_y = \mathcal{Q}^{-1}(\mathbf{z}) := \{\mathbf{y} \in \mathbb{R}^M | \mathcal{Q}(\mathbf{y}) = \mathbf{z}\}. \quad (5)$$

We introduce an auxiliary variable e denoting the quantization error:

$$e = z - y$$

with its domain

$$\mathcal{D}_e = z - \mathcal{D}_y := \{z - y | y \in \mathcal{D}_y\}. \quad (6)$$

Note that \mathcal{D}_e is unbounded as data saturation occurs.

2.1.2 1-bit CS

In the case of 1-bit quantizer we set $u_0 = -\infty$, $u_1 = 0$ and $u_2 = +\infty$. The sign information of y is preserved in the quantized measurement z . But the scaling information of y and that of x is lost. Without loss of generality, we let the 1-bit quantizer

$$\mathcal{Q}(v) = \varsigma \operatorname{sgn}(v)$$

for a scalar $v \in \mathbb{R}$ with $\varsigma \rightarrow 0_+$ (ς is an arbitrarily small positive number) and $\operatorname{sgn}(\cdot)$ being the sign function. For convenience, we set $\operatorname{sgn}(0) = 1$ (the choice is arbitrary and can be replaced by $\operatorname{sgn}(0) = -1$). Then we have $z \rightarrow 0$. To solve the signal scaling problem we impose a constraint that y has fixed unit norm, i.e.,

$$\|y\|_s = 1 \quad (7)$$

with $s \geq 1$. Different from the multi-bit quantizer case we have in such a case that

$$\mathcal{D}_y = \{y \in \mathbb{R}^M | \operatorname{sgn}(y) = \operatorname{sgn}(z), \|y\|_s = 1\}, \quad (8)$$

$$\mathcal{D}_e = \{e \in \mathbb{R}^M | \operatorname{sgn}(e) = -\operatorname{sgn}(z), \|e\|_s = 1\}. \quad (9)$$

As a result, a unified model for both multi- and 1-bit CS can be written into

$$z = Ax + e + n, \quad e \in \mathcal{D}_e, \quad (10)$$

which is the observation model to be used in this paper to recover x .

Remark 1 In the observation model (10), y is written into two variables z and e that are dependent. As an example, the statistical information of e can be obtained given that of x and n for a fixed quantizer \mathcal{Q} . For the sake of computational ease, such dependence will not be exploited in the following signal estimation process, i.e., we assume that there is no a priori information of e but its domain \mathcal{D}_e .

2.2 Bayesian Formulation of the Quantized CS Problem

In this subsection we formulate the quantized CS problem from a Bayesian perspective based on the observation model in (10). According to Remark 1 we treat e as a variable independent of x and n . The joint probability density function (PDF) $p(z, x, e)$ is decomposed as

$$p(z, x, e) = p(z|x, e) p(x) p(e).$$

We define the three distributions on the right hand side as follows.

2.2.1 Noise model

Under an assumption of white Gaussian measurement noise, i.e., $n \sim \mathcal{N}(0, \sigma^2 I)$ where σ^2 is the noise variance and I denotes an identity matrix of proper dimension, we have

$$p(z|x, e; \sigma^2) = \mathcal{N}(z|Ax + e, \sigma^2 I). \quad (11)$$

2.2.2 Sparse signal model

A sparse prior is needed for the sparse signal \mathbf{x} of interest. Here we do not give an explicit distribution to the sparse signal \mathbf{x} but denote $p(\mathbf{x})$ its PDF. Then we let $f(\mathbf{x}) = -C_1 \log p(\mathbf{x}) + C_2$ where C_1 and C_2 are proper constants. The only thing that we assume for $p(\mathbf{x})$ is that it favors entries of \mathbf{x} being zeros. As an example, a commonly used sparse prior for \mathbf{x} is a Laplace prior: $p(\mathbf{x}) = \lambda^N \exp\{-\lambda \|\mathbf{x}\|_1\}$ with λ being a positive constant. In such a case, we have $f(\mathbf{x}) = \|\mathbf{x}\|_1$ that has been extensively studied in deterministic optimization methods.

2.2.3 Quantization error model

We assume a uniform, noninformative prior for \mathbf{e} :

$$\mathbf{e} \sim U(\mathcal{D}_e) \quad (12)$$

since the only information of \mathbf{e} that we use is $\mathbf{e} \in \mathcal{D}_e$.

To obtain a maximum a posteriori (MAP) estimator of \mathbf{x} requires to integrate out \mathbf{e} from $p(\mathbf{z}, \mathbf{x}, \mathbf{e})$ that is computational intractable. We propose to estimate \mathbf{x} and \mathbf{e} simultaneously using their joint MAP estimator:

$$\begin{aligned} \{\hat{\mathbf{x}}, \hat{\mathbf{e}}\} &= \arg \max_{\mathbf{x}, \mathbf{e}} \log p(\mathbf{x}, \mathbf{e} | \mathbf{z}) \\ &= \arg \max_{\mathbf{x}, \mathbf{e}} \log p(\mathbf{z}, \mathbf{x}, \mathbf{e}) \\ &= \arg \max_{\mathbf{x}, \mathbf{e}} \log \{p(\mathbf{z} | \mathbf{x}, \mathbf{e}) p(\mathbf{x}) p(\mathbf{e})\} \\ &= \arg \min_{\mathbf{x}, \mathbf{e} \in \mathcal{D}_e} \left\{ f(\mathbf{x}) + \frac{C_1}{2\sigma^2} \|\mathbf{z} - \mathbf{e} - \mathbf{A}\mathbf{x}\|_2^2 \right\}. \end{aligned} \quad (13)$$

An equivalent form of the problem in (13) is

$$\min_{\tilde{\mathbf{x}}, \tilde{\mathbf{e}}} f(\tilde{\mathbf{x}}), \text{ subject to } \begin{cases} \|\mathbf{z} - \tilde{\mathbf{e}} - \mathbf{A}\tilde{\mathbf{x}}\|_2 \leq \epsilon, \\ \tilde{\mathbf{e}} \in \mathcal{D}_e \end{cases} \quad (14)$$

where ϵ is a proper scalar that controls the noise energy. The first constraint in (14) is to ensure the data consistency against the measurement noise. In multi-bit CS, the second one concerns data consistency due to quantization. In 1-bit CS, an additional signal scaling constraint is included in the second constraint that prevents an optimal solution for \mathbf{x} from $\mathbf{0}$. Before proceeding to our algorithm within the framework of Bayesian CS, we study in the next subsection relations of the proposed Bayesian framework with existing methods.

2.3 Relations with Existing Methods in Quantized CS

We first note that problem (14) is equivalent to the problem

$$\min_{\tilde{\mathbf{x}}, \tilde{\mathbf{y}}} f(\tilde{\mathbf{x}}), \text{ subject to } \begin{cases} \|\tilde{\mathbf{y}} - \mathbf{A}\tilde{\mathbf{x}}\|_2 \leq \epsilon, \\ \tilde{\mathbf{y}} \in \mathcal{D}_y. \end{cases} \quad (15)$$

In the following we show that many existing problem formulations of quantized CS are special cases of or related to (15), with the focus on the 1-bit case.

2.3.1 Multi-bit CS

We consider the case of ℓ_1 optimization where $f(\mathbf{x}) = \|\mathbf{x}\|_1$. In the noise free case where $\epsilon = 0$, the problem in (15) can be written into

$$\min_{\tilde{\mathbf{x}}} \|\tilde{\mathbf{x}}\|_1, \text{ subject to } \mathbf{A}\tilde{\mathbf{x}} \in \mathcal{D}_y, \quad (16)$$

which has been studied in [26]. Further, by assuming that \mathcal{Q} is a uniform unsaturated quantizer the above problem can be written into

$$\min_{\tilde{\mathbf{x}}} \|\tilde{\mathbf{x}}\|_1, \text{ subject to } \|\mathbf{z} - \mathbf{A}\tilde{\mathbf{x}}\|_\infty \leq \frac{r}{2} \quad (17)$$

that is studied in [15, 16] with r denoting the quantization bin width. While existing methods that take into account measurement noise typically mix up the quantization error and the noise, e.g., in [17], problem (15) extends existing noise free formulations to the noisy case by dealing with the two uncertainties separately.

Remark 2 *Under the assumption that all quantization errors are independent and uniformly distributed in a common interval $[-\frac{r}{2}, \frac{r}{2}]$, it is shown in [16] that the ℓ_∞ norm in problem (17) is not the best choice for the signal recovery. But it is unclear whether the result in [16] can be extended to the case of a general quantizer where the above assumption fails. It is noted that our problem formulation does not require this assumption and applies to an arbitrary quantizer. That is, by losing some optimality, we have obtained the universality.*

2.3.2 1-bit CS

In 1-bit CS problem (15) becomes

$$\min_{\tilde{\mathbf{x}}, \tilde{\mathbf{y}}} f(\tilde{\mathbf{x}}), \text{ subject to } \begin{cases} \|\tilde{\mathbf{y}} - \mathbf{A}\tilde{\mathbf{x}}\|_2 \leq \epsilon, \\ \text{sgn}(\tilde{\mathbf{y}}) = \text{sgn}(\mathbf{z}), \\ \|\tilde{\mathbf{y}}\|_s = 1. \end{cases} \quad (18)$$

In the noise free case, it can be written into

$$\min_{\tilde{\mathbf{x}}} f(\tilde{\mathbf{x}}), \text{ subject to } \begin{cases} \text{sgn}(\mathbf{A}\tilde{\mathbf{x}}) = \text{sgn}(\mathbf{z}), \\ \|\mathbf{A}\tilde{\mathbf{x}}\|_s = 1. \end{cases} \quad (19)$$

This problem with the settings $f(\mathbf{x}) = \|\mathbf{x}\|_1$ and $s = 1$ can be shown to be convex and has been studied in [24]. So by problem (18) we extend (19) to the noisy case while the authors of [24] state in [25] that “it was unclear how to modify the above convex program to account for possible noise.”

Since the third constraint in (18) serves to prevent the optimal solution for \mathbf{x} from $\mathbf{0}$, it is natural to replace it by $\|\tilde{\mathbf{x}}\|_2 = 1$, leading to the problem

$$\min_{\tilde{\mathbf{x}}, \tilde{\mathbf{y}}} f(\tilde{\mathbf{x}}), \text{ subject to } \begin{cases} \|\tilde{\mathbf{y}} - \mathbf{A}\tilde{\mathbf{x}}\|_2 \leq \epsilon, \\ \text{sgn}(\tilde{\mathbf{y}}) = \text{sgn}(\mathbf{z}), \\ \|\tilde{\mathbf{x}}\|_2 = 1. \end{cases} \quad (20)$$

In the noise free case it becomes

$$\min_{\tilde{\mathbf{x}}} f(\tilde{\mathbf{x}}), \text{ subject to } \begin{cases} \text{sgn}(\mathbf{A}\tilde{\mathbf{x}}) = \text{sgn}(\mathbf{z}), \\ \|\tilde{\mathbf{x}}\|_2 = 1, \end{cases} \quad (21)$$

which with $f(\mathbf{x}) = \|\mathbf{x}\|_1$ is the earliest formulation of the 1-bit CS problem introduced in [19] and solved respectively using RFPI in [19] and RSS in [22]. Note that for any optimal solution $(\mathbf{x}^*, \mathbf{y}^*)$ to (20), $(\mathbf{x}^*, \mathbf{y}')$ with $\mathbf{y}' = \text{sgn}(\mathbf{z}) \odot (\text{sgn}(\mathbf{z}) \odot \mathbf{A}\mathbf{x}^*)_+$ is also its optimal solution by

$$\|\mathbf{y}' - \mathbf{A}\mathbf{x}^*\|_2 \leq \|\mathbf{y}^* - \mathbf{A}\mathbf{x}^*\|_2 \leq \epsilon,$$

where $(v)_+ = \max\{v, 0\}$ for a scalar v and operates elementwise for a vector. So we can set $\tilde{\mathbf{y}} = \text{sgn}(\mathbf{z}) \odot (\text{sgn}(\mathbf{z}) \odot \mathbf{A}\tilde{\mathbf{x}})_+$ in (20), leading to that problem (20) is equivalent to the problem

$$\min_{\tilde{\mathbf{x}}} f(\tilde{\mathbf{x}}), \text{ subject to } \begin{cases} \|(\text{sgn}(\mathbf{z}) \odot \mathbf{A}\tilde{\mathbf{x}})_-\|_2 \leq \epsilon, \\ \|\tilde{\mathbf{x}}\|_2 = 1 \end{cases} \quad (22)$$

where $(v)_- = \min\{v, 0\}$ for a scalar v and operates elementwise for a vector.

The noise information is used in above formulations that account for the measurement noise. If it is unknown but the signal sparsity information is known, (22) can be written into

$$\min_{\tilde{\mathbf{x}}} \|(\text{sgn}(\mathbf{z}) \odot \mathbf{A}\tilde{\mathbf{x}})_-\|_2, \text{ subject to } \begin{cases} f(\tilde{\mathbf{x}}) \leq S, \\ \|\tilde{\mathbf{x}}\|_2 = 1, \end{cases} \quad (23)$$

where the constant S refers to the sparsity information. BIHT- ℓ_2 is proposed in [23] to solve (23) with $f(\mathbf{x}) = \|\mathbf{x}\|_0$. Moreover, the objective function in (23) has been also studied in [19,21,27]. By replacing the ℓ_2 norm in the objective function in (23) by ℓ_1 norm (that corresponds to assuming a Laplace distributed noise rather than Gaussian) it gives:

$$\min_{\tilde{\mathbf{x}}} \|(\text{sgn}(\mathbf{z}) \odot \mathbf{A}\tilde{\mathbf{x}})_-\|_1, \text{ subject to } \begin{cases} f(\tilde{\mathbf{x}}) \leq S, \\ \|\tilde{\mathbf{x}}\|_2 = 1, \end{cases} \quad (24)$$

which with $f(\mathbf{x}) = \|\mathbf{x}\|_0$ is solved using BIHT in [23]. Problem (24) searches for a solution in a sparse region (defined by the feasible domain) that maximizes the data consistency by minimizing the sum of the negative part of $\text{sgn}(\mathbf{z}) \odot \mathbf{A}\tilde{\mathbf{x}}$. A related (not equivalent) problem is thus to maximize the sum of the positive part of $\text{sgn}(\mathbf{z}) \odot \mathbf{A}\tilde{\mathbf{x}}$, i.e., to maximize $\|(\text{sgn}(\mathbf{z}) \odot \mathbf{A}\tilde{\mathbf{x}})_+\|_1$, subject to the same constraints. By combining the two objective functions, i.e., to maximize $\|(\text{sgn}(\mathbf{z}) \odot \mathbf{A}\tilde{\mathbf{x}})_+\|_1 - \|(\text{sgn}(\mathbf{z}) \odot \mathbf{A}\tilde{\mathbf{x}})_-\|_1 = \text{sgn}^T(\mathbf{z}) \mathbf{A}\tilde{\mathbf{x}}$, it gives the problem (with $f(\mathbf{x}) = \|\mathbf{x}\|_1$)

$$\max_{\tilde{\mathbf{x}}} \text{sgn}^T(\mathbf{z}) \mathbf{A}\tilde{\mathbf{x}}, \text{ subject to } \begin{cases} \|\tilde{\mathbf{x}}\|_1 \leq S, \\ \|\tilde{\mathbf{x}}\|_2 = 1, \end{cases} \quad (25)$$

whose convex relaxation given by

$$\max_{\tilde{\mathbf{x}}} \text{sgn}^T(\mathbf{z}) \mathbf{A}\tilde{\mathbf{x}}, \text{ subject to } \begin{cases} \|\tilde{\mathbf{x}}\|_1 \leq S, \\ \|\tilde{\mathbf{x}}\|_2 \leq 1 \end{cases} \quad (26)$$

has been studied in [25] and shown to give guaranteed signal recovery accuracy under mild conditions. We note that it is easy to show that problem (25) has the same guarantee as (26) following the analysis in [25].

Remark 3

- (1) Problem (18) with $f(\mathbf{x}) = \|\mathbf{x}\|_1$ and $s = 1$, and problem (26) are two parallel convex formulations of noisy 1-bit CS. Problem (18) requires the noise information but not the signal sparsity level while it is an opposite situation for problem (26). Since guaranteed signal recovery performance has been proven for (26), we believe it is true that a similar result holds for (18), which we pose as an open problem and is beyond the scope of this paper.
- (2) By (22) we see that the effective noise is $(\text{sgn}(\mathbf{z}) \odot \mathbf{A}\mathbf{x})_-$ in 1-bit CS where by “effective noise” we refer to a noise that has the minimum energy and leads to the same measurement. So, the effective noise is generally a sparse noise whose sparsity level depends on the sign flip rate of the measurement. It explains why the noise can be modeled by a Laplace distribution as in (24). Moreover, it explains why BIHT outperforms BIHT- ℓ_2 in the high SNR region, as reported in [23], where the effective noise is very sparse.
- (3) It can be shown that the energy of the effective noise $(\text{sgn}(\mathbf{z}) \odot \mathbf{A}\mathbf{x})_-$ is much smaller than that of the true noise $\mathbf{y} - \mathbf{A}\mathbf{x}$. From this point of view, we may say that 1-bit CS is robust to the measurement noise.




Figure 1: Directed graphical model that encodes the joint PDF in (30) of the Bayesian model. Nodes denoted with circles correspond to random variables, while nodes denoted with squares correspond to parameters of the model. Doubly circled z is the observation while single circled nodes represent hidden variables.

3 Q-VMP: Variational Message Passing for Quantized CS

3.1 Model Selection

We assume that the noise variance σ^2 is known. Though it can be estimated by assuming an inverse Gamma prior for it in the case where it is unknown as in [6, 28], its estimate is inaccurate due to an “identifiability issue” as addressed in [29]. For the sparse signal \mathbf{x} , we adopt a three-layer, Gaussian-Gamma-Gamma, hierarchical prior introduced in [9]:

$$p(\mathbf{x}; \epsilon, c, d) = \iint p(\mathbf{x}|\alpha) p(\alpha|\eta; \epsilon) p(\eta; c, d) d\alpha d\eta$$

where

$$p(\mathbf{x}|\alpha) = \mathcal{N}(\mathbf{x}|\mathbf{0}, \mathbf{\Lambda}), \quad (27)$$

$$p(\alpha|\eta; \epsilon) = \prod_{i=1}^N \Gamma(\alpha_i|\epsilon, \eta), \quad (28)$$

$$p(\eta; c, d) = \Gamma(\eta|c, d) \quad (29)$$

with $\mathbf{\Lambda} = \text{diag}(\alpha)$ and constants ϵ, c, d . For a Gamma distributed variable $u \sim \Gamma(c, d)$, its PDF is $\Gamma(u|c, d) = \frac{d^c}{\Gamma(c)} u^{c-1} \exp(-du)$ with $\Gamma(c)$ being the Gamma function. By [9] the constants ϵ, c, d satisfy that $0 \leq \epsilon \leq 1$, $c, d \geq 0$. In this paper, we adopt $c = 1, d = 0$ to make the prior for η in (29) noninformative (flat on \mathbb{R}_+). Further, we choose $\epsilon = 0$ since a smaller ϵ leads to a sparser prior and an estimator that approximates a hard-thresholding rule according to [9]. Readers are referred to [9] for more properties of the Gaussian-Gamma-Gamma prior and its relations with other sparse estimation techniques.

In 1-bit CS, we let \mathbf{y} have unit ℓ_2 norm in (7), leading to that $\|\mathbf{e}\|_2 = 1$ in (9). As a result, we have the joint PDF of the observation model (10):

$$p(\mathbf{z}, \mathbf{x}, \mathbf{e}, \alpha, \eta) = p(\mathbf{z}|\mathbf{x}, \mathbf{e}) p(\mathbf{x}|\alpha) p(\alpha|\eta) p(\eta) p(\mathbf{e}) \quad (30)$$

with the distributions on the right hand side as defined respectively by (11), (27), (28), (29) and (12). A directed graphical model that encodes the factorization of the joint PDF in (30) is shown in Fig. 1.

3.2 Q-VMP Algorithm

It is known that Bayesian inference is based on the posterior distribution $p(\mathbf{x}, \mathbf{e}, \alpha, \eta|\mathbf{z}) = p(\mathbf{z}, \mathbf{x}, \mathbf{e}, \alpha, \eta) / p(\mathbf{z})$. However, such an exact posterior distribution is intractable since $p(\mathbf{z}) = \int \cdots \int p(\mathbf{z}, \mathbf{x}, \mathbf{e}, \alpha, \eta) d\mathbf{x} d\mathbf{e} d\alpha d\eta$ cannot be expressed explicitly.

A variational inference approach [30, 31] is adopted in this paper. Denote $\mathbf{V} = \{\mathbf{x}, \mathbf{e}, \alpha, \eta\}$ the set of all unknown variables to be estimated. The goal in variational inference is to find a tractable distribution $q(\mathbf{V})$ that closely approximates the true posterior distribution $p(\mathbf{V}|\mathbf{z})$. To do this, some family of distributions that has enough flexibility is firstly chosen to represent $q(\mathbf{V})$. Then the task is to find a member of the family that minimizes the

Kullback-Leibler (KL) divergence between the true posterior $p(\mathbf{V}|\mathbf{z})$ and the variational approximation $q(\mathbf{V})$. A commonly used variational distribution $q(\mathbf{V})$ is such that disjoint groups of variables are independent, i.e., $q(\mathbf{V})$ has a factorized form $q(\mathbf{V}) = q(\mathbf{x})q(\mathbf{e})q(\boldsymbol{\alpha})q(\boldsymbol{\eta})$. Variational message passing (VMP) is proposed in [8] for the variational inference using a message passing procedure on a graphical model. In VMP, the variational distributions $q(\mathbf{x})$, $q(\mathbf{e})$, $q(\boldsymbol{\alpha})$, $q(\boldsymbol{\eta})$ are iteratively updated to monotonically decrease the KL divergence and thus has guaranteed convergence. Readers are referred to [8] for more details of VMP. The updates of $q(\mathbf{x})$, $q(\boldsymbol{\alpha})$, $q(\boldsymbol{\eta})$ are similar to those in [9] because of the similarity between quantized and conventional CS. $q(\mathbf{e})$ is given complete flexibility in multi-bit CS as $q(\mathbf{x})$, $q(\boldsymbol{\alpha})$ and $q(\boldsymbol{\eta})$. We constrain $q(\mathbf{e})$ in 1-bit CS such that

$$q(\mathbf{e}) = \delta(\mathbf{e} - \mathbf{e}^0) \quad (31)$$

due to a computational issue as to be discussed in Remark 4, where $\delta(\cdot)$ is a delta function and $\mathbf{e}^0 \in \mathbb{R}^M$ is to be estimated. Note that (31) is equivalent to the complete flexibility in the noise free case as to be illustrated in Subsection 3.3.

3.2.1 Updates of $q(\mathbf{x})$, $q(\boldsymbol{\alpha})$ and $q(\boldsymbol{\eta})$

According to [8] we have that

$$\begin{aligned} q(\mathbf{x}) &\propto \exp \left\{ \langle \ln p(\mathbf{z}|\mathbf{x}, \mathbf{e}) \rangle_{q(\mathbf{e})} \langle \ln p(\mathbf{x}|\boldsymbol{\alpha}) \rangle_{q(\boldsymbol{\alpha})} \right\} \\ &\propto \exp \left\{ -\frac{1}{2} (\mathbf{x} - \boldsymbol{\mu})^T \boldsymbol{\Sigma}^{-1} (\mathbf{x} - \boldsymbol{\mu}) \right\}, \end{aligned}$$

and thus $q(\mathbf{x})$ is a Gaussian distribution $\mathcal{N}(\mathbf{x}|\boldsymbol{\mu}, \boldsymbol{\Sigma})$ with the mean $\boldsymbol{\mu}$ and covariance $\boldsymbol{\Sigma}$:

$$\boldsymbol{\mu} = \langle \mathbf{x} \rangle_{q(\mathbf{x})} = \sigma^{-2} \boldsymbol{\Sigma} \mathbf{A}^T \left(\mathbf{z} - \langle \mathbf{e} \rangle_{q(\mathbf{e})} \right), \quad (32)$$

$$\boldsymbol{\Sigma} = \left(\sigma^{-2} \mathbf{A}^T \mathbf{A} + \langle \boldsymbol{\Lambda}^{-1} \rangle_{q(\boldsymbol{\alpha})} \right)^{-1}. \quad (33)$$

For $\boldsymbol{\alpha}$ we have

$$\begin{aligned} q(\boldsymbol{\alpha}) &\propto \exp \left\{ \langle \ln p(\mathbf{x}|\boldsymbol{\alpha}) \rangle_{q(\mathbf{x})} \langle \ln p(\boldsymbol{\alpha}|\boldsymbol{\eta}) \rangle_{q(\boldsymbol{\eta})} \right\} \\ &\propto \prod_{n=1}^N \alpha_n^{\epsilon - \frac{3}{2}} \exp \left\{ -\frac{1}{2} \alpha_n^{-1} \langle x_n^2 \rangle_{q(\mathbf{x})} - \alpha_n \langle \eta \rangle_{q(\boldsymbol{\eta})} \right\} \end{aligned}$$

where $\langle x_n^2 \rangle_{q(\mathbf{x})} = \mu_n^2 + \Sigma_{nn}$. The expression on the right hand side is the product of generalized inverse Gaussian (GIG) PDFs and thus we have for any $i \in \mathbb{R}$ [32]:

$$\langle \alpha_n^i \rangle_{q(\boldsymbol{\alpha})} = \left(\frac{\langle x_n^2 \rangle_{q(\mathbf{x})}}{2 \langle \eta \rangle_{q(\boldsymbol{\eta})}} \right)^{\frac{i}{2}} \frac{\mathcal{K}_{\epsilon+i-\frac{1}{2}} \left(\sqrt{2 \langle \eta \rangle_{q(\boldsymbol{\eta})} \langle x_n^2 \rangle_{q(\mathbf{x})}} \right)}{\mathcal{K}_{\epsilon-\frac{1}{2}} \left(\sqrt{2 \langle \eta \rangle_{q(\boldsymbol{\eta})} \langle x_n^2 \rangle_{q(\mathbf{x})}} \right)} \quad (34)$$

where $\mathcal{K}_\nu(\cdot)$ is the modified Bessel function of the second kind and order $\nu \in \mathbb{R}$. The case of $i = -1$ in (34) gives the evaluation of $\langle \boldsymbol{\Lambda}^{-1} \rangle_{q(\boldsymbol{\alpha})}$ used in (33), and the case of $i = 1$ gives the calculation of $\langle \alpha_n \rangle_{q(\boldsymbol{\alpha})}$ used in a later expression in (35). The update of $q(\boldsymbol{\eta})$ can be shown to be $q(\boldsymbol{\eta}) = \Gamma \left(\eta | N\epsilon + c, \sum_{n=1}^N \langle \alpha_n \rangle_{q(\boldsymbol{\alpha})} + d \right)$. The first moment of η used in (34) is given as

$$\langle \eta \rangle_{q(\boldsymbol{\eta})} = \frac{N\epsilon + c}{\sum_{n=1}^N \langle \alpha_n \rangle_{q(\boldsymbol{\alpha})} + d}. \quad (35)$$

3.2.2 Update of $q(e)$ in multi-bit CS

In multi-bit CS we have

$$\begin{aligned} q(e) &\propto \exp \left\{ \langle \ln p(z|x, e) \rangle_{q(x)} \right\} p(e) \\ &\propto \exp \left\{ -\frac{1}{2} \sigma^{-2} \left\langle \|z - e - \mathbf{A}x\|_2^2 \right\rangle_{q(x)} \right\} I_e(\mathcal{D}_e) \\ &\propto \exp \left\{ -\frac{1}{2} \sigma^{-2} \|e - (z - \mathbf{A}\mu)\|_2^2 \right\} I_e(\mathcal{D}_e) \end{aligned} \quad (36)$$

where $I_e(\mathcal{D}_e)$ is an indicator function that equals to 1 if $e \in \mathcal{D}_e$ or 0 otherwise. Hence, $q(e)$ is the product of PDFs of truncated Gaussian distributions, i.e., for each $m = 1, \dots, M$, $q(e_m)$ is the PDF of a truncated Gaussian distribution. As a result, the first moment of e_m , $m = 1, \dots, M$, used in (32) can be given in closed form after some derivations using the PDF $\phi(\cdot)$ and cumulative distribution function (CDF) $\Phi(\cdot)$ of a standard Gaussian distribution:

$$\langle e_m \rangle_{q(e)} = \sigma \frac{\phi(l_{e_m}) - \phi(u_{e_m})}{\Phi(u_{e_m}) - \Phi(l_{e_m})} + \mu_{e_m}, \quad (37)$$

where $\mu_{e_m} = (z - \mathbf{A}\mu)_m$, l_{e_m} and u_{e_m} satisfy that $\mathcal{D}_{e_m} = [\sigma l_{e_m} + \mu_{e_m}, \sigma u_{e_m} + \mu_{e_m}]$ with \mathcal{D}_{e_m} denoting the domain of e_m , $\phi(u) = \frac{1}{\sqrt{2\pi}} \exp\left\{-\frac{u^2}{2}\right\}$ and $\Phi(u) = \int_{-\infty}^u \phi(t) dt$ for $u \in \mathbb{R}$.

Remark 4 Consider the case where $q(e)$ is given the complete flexibility in 1-bit CS. Note that entries of a point in \mathcal{D}_e are no longer independent of each other in such a case, leading to that $q(e)$ is the PDF of a truncated multi-variable Gaussian distribution with e constrained in a nonconvex set \mathcal{D}_e defined in (9). As a result, the calculation of $\langle e \rangle_{q(e)}$ is in general computationally intractable in our considered CS problems where the dimension of e is large.

3.2.3 Update of $q(e)$ in 1-bit CS

According to [8], this is equivalent to finding an MAP estimator of e with its posterior distribution defined in (36). So we have

$$\langle e \rangle_{q(e)} = e^0 \quad (38)$$

with

$$\begin{aligned} e^0 &= \arg \max_{e \in \mathbb{R}^M} \left\{ \exp \left\{ -\frac{1}{2} \sigma^{-2} \|e - (z - \mathbf{A}\mu)\|_2^2 \right\} I_e(\mathcal{D}_e) \right\} \\ &= \arg \min_{e \in \mathcal{D}_e} \|e - (z - \mathbf{A}\mu)\|_2 \\ &= \mathcal{P}_{\mathcal{D}_e}(z - \mathbf{A}\mu) \end{aligned}$$

where \mathcal{D}_e is defined in (9) and $\mathcal{P}_{\mathcal{D}}(v)$ denotes a projection of a point v onto a set \mathcal{D} . The calculation of $\mathcal{P}_{\mathcal{D}_e}(\cdot)$ with the nonconvex set \mathcal{D}_e is provided in the following lemma.

Lemma 1 For a vector $v \in \mathbb{R}^M$, let $\bar{v} = -\text{sgn}(z) \odot v$. Denote \mathcal{I} the index set of all positive entries of \bar{v} . Let \mathcal{I}^c be its complementary set. If \mathcal{I} is nonempty, then let $e^* \in \mathbb{R}^M$ with $e_{\mathcal{I}}^* = \frac{v_{\mathcal{I}}}{\|v_{\mathcal{I}}\|_2}$ and $e_{\mathcal{I}^c}^* = \mathbf{0}$. Otherwise, let $e_i^* = \begin{cases} -\text{sgn}(z_{i_0}), & \text{if } i = i_0; \\ 0, & \text{otherwise,} \end{cases}$ for $i = 1, \dots, M$, where $i_0 \in \{1, \dots, M\}$ with $\bar{v}_{i_0} = \max(\bar{v})$. Then we have $e^* = \mathcal{P}_{\mathcal{D}_e}(v)$ with \mathcal{D}_e as defined in (9).




Figure 2: An illustration of Lemma 1 with nonnegative entries of e . The unit circle in the first quadrant composes of \mathcal{D}_e . Projections of four possible v 's are shown.

Proof See Appendix.

Lemma 1 tells how to calculate the projection onto the nonconvex set \mathcal{D}_e defined in (9). An illustration of Lemma 1 is presented in Fig. 2, where we consider the two dimensional case with both entries of e nonnegative. The unit circle in the first quadrant composes of \mathcal{D}_e . Projections of four possible v 's are shown. The resulting algorithm is summarized in Algorithm 1, named as variational message passing with quantization (Q-VMP).

Algorithm 1: Q-VMP

Input: sensing matrix \mathbf{A} , quantized measurement z , domain of quantization error \mathcal{D}_e , and noise variance σ^2 .

1. initialize $\langle \alpha_n^{-1} \rangle_{q(\alpha)}$, $n = 1, \dots, N$, $\langle \eta \rangle_{q(\eta)}$ and $\langle e \rangle_{q(e)}$;
2. **while** not converged **do**
3. update Σ according to (33);
4. update μ according to (32);
5. update $\langle \alpha_n^{-1} \rangle_{q(\alpha)}$ and $\langle \alpha_n \rangle_{q(\alpha)}$, $n = 1, \dots, N$, according to (34);
6. update $\langle \eta \rangle_{q(\eta)}$ according to (35);
7. update $\langle e \rangle_{q(e)}$ according to (37) in multi-bit CS and (38) in 1-bit CS respectively;
8. **end while**

Output: recovered signal $\hat{x} = \mu$.

3.3 The Noise Free Case

In this subsection we consider Q-VMP in the noise free case. We first consider the data consistency. A consistent recovery means that the observation can be reproduced from the recovered signal. Empirical results suggest that a consistent recovery result in less errors [16, 22]. A theoretical proof is provided in [23] on the 1-bit case. The following analysis applies to both multi- and 1-bit CS. Taking $\sigma^2 \rightarrow 0$ at both sides of (32) and (33) gives

$$\begin{aligned} \mu &\rightarrow \langle \Lambda^{-1} \rangle_{q(\alpha)}^{-\frac{1}{2}} \left(\mathbf{A} \langle \Lambda^{-1} \rangle_{q(\alpha)}^{-\frac{1}{2}} \right)^\dagger (z - e), \\ \Sigma &\rightarrow \langle \Lambda^{-1} \rangle_{q(\alpha)}^{-1} - \langle \Lambda^{-1} \rangle_{q(\alpha)}^{-\frac{1}{2}} \left(\mathbf{A} \langle \Lambda^{-1} \rangle_{q(\alpha)}^{-\frac{1}{2}} \right)^\dagger \mathbf{A} \langle \Lambda^{-1} \rangle_{q(\alpha)}^{-1}. \end{aligned}$$

Thus we have

$$\mathbf{A}\mu \rightarrow z - e \in \mathcal{D}_y,$$

i.e., $\mathcal{Q}(\mathbf{A}\boldsymbol{\mu}) \rightarrow \mathbf{z}$, which indicates that the recovered signal reproduces the observation at each iteration.

We next consider the update of $q(\mathbf{e})$ in such a case. As $\sigma^2 \rightarrow 0$ we see that $q(\mathbf{e})$ degenerates into a single-point distribution by (36), that coincides with the stricter assumption in (31) in 1-bit CS.

3.4 Pruning a Basis Function

The most difficult computation of Q-VMP is the calculation of $\boldsymbol{\Sigma}$ that is the inverse of an $N \times N$ matrix. Using the Woodbury matrix identity, we have

$$\boldsymbol{\Sigma} = \langle \boldsymbol{\Lambda}^{-1} \rangle_{q(\boldsymbol{\alpha})}^{-1} - \langle \boldsymbol{\Lambda}^{-1} \rangle_{q(\boldsymbol{\alpha})}^{-1} \mathbf{A}^T \mathbf{C}^{-1} \mathbf{A} \langle \boldsymbol{\Lambda}^{-1} \rangle_{q(\boldsymbol{\alpha})}^{-1}$$

with $\mathbf{C} = \sigma^2 \mathbf{I} + \mathbf{A} \langle \boldsymbol{\Lambda}^{-1} \rangle_{q(\boldsymbol{\alpha})}^{-1} \mathbf{A}^T$ being an $M \times M$ matrix. Hence, to calculate $\boldsymbol{\Sigma}$ needs $O(\min\{N^3, N^2M\})$ operations. It is noted that if Q-VMP produces some $\langle \alpha_n^{-1} \rangle_{q(\boldsymbol{\alpha})} \rightarrow +\infty$ with $n \in \{1, \dots, N\}$, then the corresponding basis \mathbf{A}_n can be removed from the model. To further speed up Q-VMP, we prune a basis \mathbf{A}_n from the model (to reduce N) when the corresponding parameter $\langle \alpha_n^{-1} \rangle_{q(\boldsymbol{\alpha})}$ is larger than a certain threshold $\tau_{pruning}$. Similar basis pruning approaches have been used in [4, 9].

4 Numerical Simulations

In this section, we study the performance of the proposed observation model and Q-VMP algorithm in comparison with existing ones by numerical simulations.

4.1 Experimental Setup

Model Parameters in Q-VMP: We set $\epsilon = 0$, $c = 1$ and $d = 0$ in the Gaussian-Gamma-Gamma prior as discussed in Subsection 2.2.

Quantizer: In multi-bit CS, a uniform unsaturated quantizer is defined in (4) with $L = 2^B$, equispaced u_0, u_1, \dots, u_L and $v_i = (u_i + u_{i+1})/2$, $i = 1, \dots, L-1$. In addition, we let $u_L = \|\mathbf{y}\|_\infty$ and $u_0 = -\|\mathbf{y}\|_\infty$ in each trial. For a saturated quantizer, we set $u_0 = -\infty$, $u_L = +\infty$.

CS problem generation: In our experiment, we set $N = 500$, $K = 10$, and vary the bit budget (total bits of all measurements) in $\{50, 100, \dots, 1000\}$. In each trial, a K -sparse signal of length N is generated with Gaussian distributed nonzero entries and then scaled to unit norm. Entries of the sensing matrix \mathbf{A} are generated independently according to a Gaussian distribution $\mathcal{N}(0, M^{-1})$. Thus the noise free measurement $\mathbf{y}^0 = \mathbf{A}\mathbf{x}$ has unit norm in expectation. To obtain a desired SNR, a white Gaussian measurement noise \mathbf{n} is added with the noise variance $\sigma^2 = M^{-1}10^{-\frac{\text{SNR}}{10}}$. Then the quantized measurement $\mathbf{z} = \mathcal{Q}(\mathbf{y})$ is preserved for the following signal recovery.

Performance metrics: Three metrics are considered, including reconstruction SNR (RSNR), sparsity level of the recovered signal and computational speed. RSNR is defined as

$$\text{RSNR} = -20 \log_{10} \|\mathbf{x} - \hat{\mathbf{x}}\|_2$$

where $\hat{\mathbf{x}}$ denotes the recovered signal of \mathbf{x} . The sparsity level is measured by the support size of the recovered signal. The computational speed is measured by the CPU time usage. All results are averaged over 200 trials.

4.2 Model Efficiency

We first study the efficiency of the observation model in (10) introduced in this paper for quantized CS. We consider the multi-bit CS problem with a uniform unsaturated quantizer as an example. In existing methods that account for measurement noise, e.g., in [17], the quantization error and the noise are typically coupled and treated as a Gaussian




Figure 3: Reconstruction SNRs of VMP algorithms implemented respectively based on the proposed observation model in (10), an existing one that couples the quantization error and measurement noise, and conventional CS (oracle-aided quantized CS) as an upper boundary.

noise (only the energy information is used). Then the quantized CS problem is transformed into a conventional one. We refer to this formulation as existing method hereafter. In this subsection we compare the signal recovery performance of the proposed formulation in (10) with the existing one. Naturally, we use the proposed Q-VMP algorithm with our formulation. A corresponding algorithm for the existing formulation is thus VMP introduced in [9] for conventional CS. The latter algorithm can be considered as a simplified version of Q-VMP with the quantization error e fixed throughout the algorithm. In addition, we consider the performance of conventional CS for comparison where the true-valued measurements are used. The conventional CS problem can be considered as one in quantized CS by using oracle information of the quantization error. So its performance acts as an upper boundary of the quantized CS problem.

In our experiment, we set $\text{SNR} = 30\text{dB}$ and the bit depth $B = 4$ that leads to the number of quantized measurements varying from 12 to 250. Q-VMP is implemented as follows.

Q-VMP: We initialize $\langle \alpha_n^{-1} \rangle_{q(\alpha)} = 1/|\mathbf{A}_n^T \mathbf{z}|$, $n = 1, \dots, N$, $\langle \eta \rangle_{q(\eta)} = 1$ and $\langle e \rangle_{q(e)} = \mathbf{0}$. We set $\tau_{\text{pruning}} = 10^4$. Q-VMP is terminated if $\frac{\|\tilde{\alpha}^j - \tilde{\alpha}^{j-1}\|_2}{\|\tilde{\alpha}^{j-1}\|_2} < 10^{-5}$ or the maximum number of iterations, set to 2000, is reached,

where $\tilde{\alpha} = [\langle \alpha_1^{-1} \rangle_{q(\alpha)}^{-1}, \dots, \langle \alpha_N^{-1} \rangle_{q(\alpha)}^{-1}]^T$ and the superscript j indicates the iteration.

The VMP algorithm for the other two cases is similarly implemented. The true noise variance is used in Q-VMP and conventional CS. For VMP with the existing formulation, we set the noise variance to $r^2/12 + \sigma^2$ where $r = u_1 - u_0$ denotes the quantization bin width. This noise variance corresponds to a Gaussian noise whose energy is comparable with that of $e + \mathbf{n}$ under the assumption that e is uniformly distributed and independent of \mathbf{n} . Reconstruction SNRs of the three methods are depicted in Fig. 3. It can be seen that the VMP algorithm based on the proposed observation model is consistently better than that with the existing formulation though there is still a large gap between it and the oracle-aided case.




Figure 4: Performance comparison of Q-VMP, BPDN and L1RML with bit depth $B = 4$. SNR = 30dB. (a) Averaged reconstruction SNR; (b) Averaged support size of recovered signal; (c) Averaged CPU time.

4.3 Performance Comparison in Multi-bit CS

4.3.1 Unsaturated quantizer

In multi-bit CS, we first consider the case of a uniform unsaturated quantizer. As in the last subsection, we set SNR = 30dB and $B = 4$. Besides Q-VMP, we also use BPDN [3] and L1RML [18] to recover the signal for comparison. Q-VMP is implemented as in the last subsection. The other two are implemented as follows.

BPDN: BPDN solves the problem in (2) with \mathbf{y} replaced by \mathbf{z} and is implemented using ℓ_1 -magic [33]. We let $\epsilon = \|\mathbf{z} - \mathbf{A}\mathbf{x}\|_2$ in our implementation to achieve the best result though it is unavailable in practice.

L1RML: L1RML solves the problem

$$\min_{\tilde{\mathbf{x}}} \{ \lambda \|\tilde{\mathbf{x}}\|_1 - \log f_{ML}(\mathbf{A}\tilde{\mathbf{x}}) \}$$

where $f_{ML}(\cdot)$ is the likelihood function of the observation with $-\log f_{ML}(\mathbf{A}\tilde{\mathbf{x}})$ being a convex function of $\tilde{\mathbf{x}}$, and $\lambda > 0$ is a regularization parameter. In general, a larger λ leads to a sparser solution. Since there are no available guidelines for the choice of λ so far, we choose λ such that the recovered signal has the optimal RSNR. Additionally, we set $\tau = \frac{\sigma^2}{\|\mathbf{A}\|_2^2}$, $\epsilon = 10^{-4}$ and $\beta = 0.5$. Readers are referred to [18] for their interpretations.

The experimental results are shown in Fig. 4, where red solid lines denote Q-VMP, black dashed dot lines denote BPDN, and blue dashed lines denote L1RML. Fig. 4(a) depicts the averaged reconstruction SNRs of the three algorithms. A significant improvement of the reconstruction SNR can be observed using the proposed Q-VMP. It is over 6dB in comparison with L1RML and about an amplitude for BPDN. Moreover, Fig. 4(b) shows that Q-VMP produces the sparsest solution. It is noted that L1RML can produce a sparser solution by setting a larger value of λ as in [18] but at the cost of a lower RSNR since it produces the optimal RSNR in our implementation. Fig. 4(c) shows that the speed of Q-VMP is comparable with that of BPDN and L1RML. Implemented with the basis pruning approach, Q-VMP is faster when more measurements are acquired since it is observed in such a case that the basis pruning approach works more efficiently.

4.3.2 Saturated Quantizer

We next consider the case of a saturated quantizer. We adopt the same experimental setup but a saturated quantizer where a noisy measurement falls in each quantization interval with the same probability. Since both the sensing matrix and measurement noise are Gaussian in the experiment, the noisy measurements are i.i.d. Gaussian $\mathcal{N}(0, M^{-1} + \sigma^2)$. Then it is easy to get the quantizer. As a result, 12.5% of the measurements are saturated in




Figure 5: Reconstruction SNRs of Q-VMP and L1RML with a saturated quantizer, as well as those with the unsaturated quantizer in Fig. 4.




Figure 6: Performance comparison of Q-VMP, BIHT and CVXP in 1-bit CS. SNR = 10dB. (a) Averaged reconstruction SNR; (b) Averaged support size of recovered signal; (c) Averaged CPU time.

expectation. BPDN is inappropriate in such a case. We compare Q-VMP only with L1RML. The averaged reconstruction SNRs of Q-VMP and L1RML are presented in Fig. 5 (red solid lines). Q-VMP obtains a RSNR of about 10dB higher than L1RML when sufficient measurements are acquired. The performance of the two algorithms on support size and speed is similar to that in the uniform quantizer case and is omitted.

The experiment above may shed light on the optimal quantizer design for Q-VMP. By comparing the performance of Q-VMP in the two quantizer scenarios, it can be seen from Fig. 5 that the saturated quantizer outperforms the uniform unsaturated one when more measurements are taken for Q-VMP while it is not so clear for L1RML. We pose the problem of the optimal quantizer design for Q-VMP as a future work.

4.4 Performance Comparison in 1-bit CS

The bit-depth $B = 1$ in 1-bit CS. We set SNR = 10dB. In such a case, 9.75% measurements flip their signs due to the noise in expectation. We compare Q-VMP with the state-of-the-art algorithms BIHT [23] and the convex programming approach in [25], denoted by CVXP. The three algorithms are implemented as follows.

Q-VMP: We initialize $\langle \alpha_n^{-1} \rangle_{q(\alpha)} = \sqrt{M} / |A_n^T \text{sgn}(z)|$, $n = 1, \dots, N$, $\langle \eta \rangle_{q(\eta)} = 1$ and $\langle e \rangle_{q(e)} = -\text{sgn}(z) / \sqrt{M}$.

As addressed in Remark 3, the effective noise level in 1-bit CS is much lower than the true one. We empirically find that it is a good choice to set the noise variance in Q-VMP to $10^{-3}\sigma^2$ (Q-VMP is slow in the case of a very small noise variance, which is the reason why we consider a lower SNR in 1-bit CS). We set $\tau_{pruning} = 10^4$ and terminate Q-VMP as in multi-bit CS. The recovered signal is finally scaled to unit norm for comparison with the original one.

BIHT: The oracle information of K is used in the hard thresholding operation, i.e., BIHT is certain to return a reconstruction with K nonzero entries. BIHT is terminated if the Hamming error of the current recovery is below the expected Hamming error or the maximum number of iterations, set to 1000, is reached. Readers are referred to [23] for the definition of the Hamming error.

CVXP: CVXP solves the convex problem in (26). It does not require the noise information but needs to know the signal sparsity level. We set $S = \|x\|_1$ in (26) (using the oracle information of $\|x\|_1$) to achieve the best result. We implement it using CVX [34] (its speed can be accelerated using faster solvers).

Our experimental results are presented in Fig. 6. In all figures, red solid lines denote Q-VMP, black dashed dot lines denote BIHT, and blue dashed lines denote CVXP. It is shown in Fig. 6(a) that the proposed Q-VMP outperforms consistently the other two algorithms in the recovery accuracy. From Fig. 6(b), it can be seen that Q-VMP produces a sparser solution than CVXP on average while BIHT uses this oracle information. Fig. 6(c) shows that the computational speed is an disadvantage of Q-VMP.

5 Conclusion and Future Work

The sparse signal recovery problem from noisy quantized compressive measurements was studied in this paper. A unified framework was presented that applies to both multi- and 1-bit CS. Moreover, the unified framework incorporates the unsaturated and saturated cases in the multi-bit quantization. Relations between the proposed unified framework and existing quantized CS methods were studied and a unified algorithm for quantized CS was proposed based on variational Bayesian inference. Numerical simulations were carried out, showing that the proposed framework and algorithm improve the signal recovery accuracy in comparison with the existing results.

A convex formulation of the noisy 1-bit CS problem has been studied in [25] with guaranteed signal recovery performance. This paper has introduced a different convex formulation (problem (18) with $f(x) = \|x\|_1$ and $s = 1$) that explicitly exploits the noise information and does not need the knowledge of the signal sparsity. One future work is to explore its theoretical guarantee. One drawback of Q-VMP is its high computational complexity due to an inversion of a high dimensional matrix at each iteration though it has been greatly alleviated with a basis pruning approach adopted in this paper. Q-VMP may suffer from computational issues in the case of very high dimensional problems. Thus another future work is to develop fast alternatives to the current implementation. Since the signal recovery accuracy in multi-bit CS is very different when a different quantizer is adopted, as shown in the present paper and in [17], to design the optimal quantizer that minimizes the signal recovery error is another interesting future research topic.

Appendix

Proof of Lemma 1

To prove Lemma 1, we first show the following three lemmas.

Lemma 2 Assume $w \in \mathbb{R}^m$, $w \succcurlyeq 0$ and $w \neq 0$. Then $v^* = \frac{C}{\|w\|_2} w$ is an optimal solution to the problem

$$\max_{v \in \mathbb{R}^m} f(v) = w^T v, \text{ subject to } \begin{cases} \|v\|_2 = C, \\ v \succcurlyeq 0 \end{cases} \quad (39)$$

with scalar $C > 0$. Consequently, $f^* = f(v^*) = C \|w\|_2$.

Proof By the Cauchy inequality we have $f(\mathbf{v}) \leq \|\mathbf{w}\|_2 \|\mathbf{v}\|_2 = C \|\mathbf{w}\|_2$. The equality holds if \mathbf{v} is in the form of \mathbf{v}^* . So, we get the conclusion.

Lemma 3 Assume $\mathbf{w} \in \mathbb{R}^m$, $\mathbf{w} \succneq \mathbf{0}$. Let $\mathbf{v}^* \in \mathbb{R}^m$ with $v_i^* = \begin{cases} C, & \text{if } i = i_0; \\ 0, & \text{otherwise,} \end{cases}$ for $i = 1, \dots, m$, where $i_0 \in \{1, \dots, n\}$ with $w_{i_0} = \min(\mathbf{w})$. Then \mathbf{v}^* is an optimal solution to the problem

$$\min_{\mathbf{v} \in \mathbb{R}^m} f(\mathbf{v}) = \mathbf{w}^T \mathbf{v}, \text{ subject to } \begin{cases} \|\mathbf{v}\|_2 = C, \\ \mathbf{v} \succneq \mathbf{0} \end{cases} \quad (40)$$

with scalar $C > 0$. Consequently, $f^* = f(\mathbf{v}^*) = C \min(\mathbf{w})$.

Proof The case of $m = 1$ is obvious. We next prove the lemma in the case of $m = 2$ and then use induction to complete the proof. Suppose $m = 2$. We need to show that if $w_1 < w_2$, then $v_1^* = C$ and $v_2^* = 0$. By substituting $v_1 = \sqrt{C^2 - v_2^2}$ into $f(\mathbf{v})$, we have

$$g(v_2) := f\left(\sqrt{C^2 - v_2^2}, v_2\right) = \sqrt{C^2 - v_2^2} w_1 + v_2 w_2. \quad (41)$$

It is easy to show that $g'(v_2) \geq 0$ if $0 \leq v_2 \leq \frac{C w_2}{\|\mathbf{w}\|_2}$, and $g'(v_2) \leq 0$ if $\frac{C w_2}{\|\mathbf{w}\|_2} \leq v_2 < C$. So, the minimum of $g(v_2)$ can only be achieved at the boundary of the interval $[0, C]$. By $g(0) = C w_1$, $g(C) = C w_2$ and $w_1 < w_2$, we have $v_2^* = 0$, $v_1^* = C$ and $f^* = C \min(\mathbf{w})$.

Suppose the lemma holds in the case $m = n - 1$ with $n > 3$. We next show that it holds in the case $m = n$. Denote $\mathbf{v}_{-1} = [v_2, \dots, v_n]^T$, $\mathbf{w}_{-1} = [w_2, \dots, w_n]^T$. Then we have $\mathbf{w}_{-1}^T \mathbf{v}_{-1} \geq \|\mathbf{v}_{-1}\|_2 \min(\mathbf{w}_{-1})$. By $\|\mathbf{v}_{-1}\|_2 = \sqrt{C^2 - v_1^2}$ we have

$$\begin{aligned} f(\mathbf{v}) &= w_1 v_1 + \mathbf{w}_{-1}^T \mathbf{v}_{-1} \\ &\geq w_1 v_1 + \min(\mathbf{w}_{-1}) \sqrt{C^2 - v_1^2} \\ &\geq C \min(\mathbf{w}). \end{aligned} \quad (42)$$

It is obvious that the equality holds if \mathbf{v} is in the form of \mathbf{v}^* .

Lemma 4 For $\mathbf{w} \in \mathbb{R}^m$, denote \mathcal{I} the index set of all positive entries of \mathbf{w} . Let \mathcal{I}^c be its complementary set. If \mathcal{I} is nonempty, then let $\mathbf{v}^* \in \mathbb{R}^m$ with $v_{\mathcal{I}}^* = \frac{\mathbf{w}_{\mathcal{I}}}{\|\mathbf{w}_{\mathcal{I}}\|_2}$ and $\mathbf{v}_{\mathcal{I}^c}^* = \mathbf{0}$. Otherwise, let $v_i^* = \begin{cases} 1, & \text{if } i = i_0; \\ 0, & \text{otherwise,} \end{cases}$ for $i = 1, \dots, m$, where $i_0 \in \{1, \dots, m\}$ with $w_{i_0} = \max(\mathbf{w})$. Then \mathbf{v}^* is an optimal solution to the problem

$$\max_{\mathbf{v} \in \mathbb{R}^m} f(\mathbf{v}) = \mathbf{w}^T \mathbf{v}, \text{ subject to } \begin{cases} \|\mathbf{v}\|_2 = 1, \\ \mathbf{v} \succneq \mathbf{0}. \end{cases} \quad (43)$$

Proof Consider first the case where \mathcal{I} is nonempty. By $\|\mathbf{v}_{\mathcal{I}}\|_2^2 + \|\mathbf{v}_{\mathcal{I}^c}\|_2^2 = \|\mathbf{v}\|_2^2 = 1$ and Lemmas 2 and 3 we have

$$\begin{aligned} f(\mathbf{v}) &= \mathbf{w}_{\mathcal{I}}^T \mathbf{v}_{\mathcal{I}} - (-\mathbf{w}_{\mathcal{I}^c}^T) \mathbf{v}_{\mathcal{I}^c} \\ &\leq \|\mathbf{v}_{\mathcal{I}}\|_2 \|\mathbf{w}_{\mathcal{I}}\|_2 - \|\mathbf{v}_{\mathcal{I}^c}\|_2 \min(-\mathbf{w}_{\mathcal{I}^c}) \\ &\leq \|\mathbf{w}_{\mathcal{I}}\|_2. \end{aligned} \quad (44)$$

It is obvious that the equality holds if \mathbf{v} is in the form of \mathbf{v}^* .

In the other case where \mathcal{I} is empty, we have $\mathbf{w} \preceq \mathbf{0}$. By $\|\mathbf{v}\|_2 = 1$ and Lemma 3 we have

$$f(\mathbf{v}) = -(-\mathbf{w}^T) \mathbf{v} \leq -\|\mathbf{v}\|_2 \min(-\mathbf{w}) = \max(\mathbf{w}). \quad (45)$$

The equality holds again if \mathbf{v} is in the form of \mathbf{v}^* .

Proof of Lemma 1: To show $e^* = \mathcal{P}_{\mathcal{D}_e}(v)$, i.e.,

$$e^* = \arg \min_{e \in \mathcal{D}_e} \|e - v\|_2,$$

we need only to show

$$e^* = \arg \max_{e \in \mathcal{D}_e} v^T e,$$

or

$$v^T e^* \geq v^T e, \text{ for any } e \in \mathcal{D}_e,$$

or further

$$\bar{v}^T (\text{sgn}(e) \odot e^*) \geq \bar{v}^T (\text{sgn}(e) \odot e), \text{ for any } e \in \mathcal{D}_e.$$

It is equivalent to showing that $-\text{sgn}(z) \odot e^*$ is an optimal solution to the problem

$$\max_{\bar{e}} \bar{v}^T \bar{e}, \text{ subject to } \begin{cases} \|\bar{e}\|_2 = 1, \\ \bar{e} \succcurlyeq \mathbf{0}. \end{cases}$$

This is a direct result by applying Lemma 4. So, we complete the proof.

References

- [1] E. Candès, “Compressive sampling,” in *Proceedings of the International Congress of Mathematicians*, vol. 3. Citeseer, 2006, pp. 1433–1452.
- [2] D. Donoho, “Compressed sensing,” *IEEE Transactions on Information Theory*, vol. 52, no. 4, pp. 1289–1306, 2006.
- [3] E. Candès, J. Romberg, and T. Tao, “Stable signal recovery from incomplete and inaccurate measurements,” *Communications on Pure and Applied Mathematics*, vol. 59, no. 8, pp. 1207–1223, 2006.
- [4] M. Tipping, “Sparse bayesian learning and the relevance vector machine,” *The Journal of Machine Learning Research*, vol. 1, pp. 211–244, 2001.
- [5] D. Wipf and B. Rao, “Sparse bayesian learning for basis selection,” *Signal Processing, IEEE Transactions on*, vol. 52, no. 8, pp. 2153–2164, 2004.
- [6] S. Babacan, R. Molina, and A. Katsaggelos, “Bayesian compressive sensing using laplace priors,” *IEEE Transactions on Image Processing*, vol. 19, no. 1, pp. 53–63, 2010.
- [7] D. MacKay, “Bayesian interpolation,” *Neural computation*, vol. 4, no. 3, pp. 415–447, 1992.
- [8] J. Winn and C. Bishop, “Variational message passing,” *Journal of Machine Learning Research*, vol. 6, no. 1, p. 661, 2006.
- [9] N. Pedersen, D. Shutin, C. Manchón, and B. Fleury, “Sparse estimation using bayesian hierarchical prior modeling for real and complex models,” *Arxiv preprint, availabel at <http://arxiv.org/abs/1108.4324>*, 2011.
- [10] L. He and L. Carin, “Exploiting structure in wavelet-based bayesian compressive sensing,” *IEEE Transactions on Signal Processing*, vol. 57, no. 9, pp. 3488–3497, 2009.
- [11] D. Donoho, A. Maleki, and A. Montanari, “The noise-sensitivity phase transition in compressed sensing,” *IEEE Transactions on Information Theory*, vol. 57, no. 10, pp. 6920–6941, 2011.

- [12] M. Stojnic, "Various thresholds for ℓ_1 -optimization in compressed sensing," *Arxiv preprint, available online at <http://arxiv.org/abs/0907.3666>*, 2009.
- [13] Z. Yang, C. Zhang, and L. Xie, "On phase transition of compressed sensing in the complex domain," *IEEE Signal Processing Letters*, vol. 19, no. 1, pp. 47–50, 2012.
- [14] A. Maleki, L. Anitori, Z. Yang, and R. Baraniuk, "Asymptotic analysis of complex lasso via complex approximate message passing (camp)," *Arxiv preprint, available online at <http://arxiv.org/abs/1108.0477>*, 2011.
- [15] E. Candès and T. Tao, "Near-optimal signal recovery from random projections: Universal encoding strategies?" *IEEE Transactions on Information Theory*, vol. 52, no. 12, pp. 5406–5425, 2006.
- [16] L. Jacques, D. Hammond, and J. Fadili, "Dequantizing compressed sensing: When oversampling and non-gaussian constraints combine," *IEEE Transactions on Information Theory*, vol. 57, no. 1, pp. 559–571, 2011.
- [17] J. Laska, P. Boufounos, M. Davenport, and R. Baraniuk, "Democracy in action: Quantization, saturation, and compressive sensing," *Applied and Computational Harmonic Analysis*, vol. 31, no. 3, pp. 429–443, 2011.
- [18] A. Zymnis, S. Boyd, and E. Candes, "Compressed sensing with quantized measurements," *Signal Processing Letters, IEEE*, vol. 17, no. 2, pp. 149–152, 2010.
- [19] P. Boufounos and R. Baraniuk, "1-bit compressive sensing," in *42nd Annual Conference on Information Sciences and Systems (CISS)*. IEEE, 2008, pp. 16–21.
- [20] J. Laska and R. Baraniuk, "Regime change: Bit-depth versus measurement-rate in compressive sensing," *Arxiv preprint, available online at <http://arxiv.org/abs/1110.3450>*, 2011.
- [21] P. Boufounos, "Greedy sparse signal reconstruction from sign measurements," in *2009 Conference Record of the Forty-Third Asilomar Conference on Signals, Systems and Computers*. IEEE, 2009, pp. 1305–1309.
- [22] J. Laska, Z. Wen, W. Yin, and R. Baraniuk, "Trust, but verify: Fast and accurate signal recovery from 1-bit compressive measurements," *IEEE Transactions on Signal Processing*, vol. 59, no. 11, pp. 5289–5301, 2011.
- [23] L. Jacques, J. Laska, P. Boufounos, and R. Baraniuk, "Robust 1-bit compressive sensing via binary stable embeddings of sparse vectors," *Arxiv preprint, available online at <http://arxiv.org/abs/1104.3160>*, 2011.
- [24] Y. Plan and R. Vershynin, "One-bit compressed sensing by linear programming," *Available online at <http://arxiv.org/abs/1109.4299>*, 2011.
- [25] —, "Robust 1-bit compressed sensing and sparse logistic regression: A convex programming approach," *Available online at <http://arxiv.org/abs/1202.1212>*, 2012.
- [26] W. Dai and O. Milenkovic, "Information theoretical and algorithmic approaches to quantized compressive sensing," *IEEE Transactions on Communications*, vol. 59, no. 7, pp. 1857–1866, 2011.
- [27] P. Boufounos, "Reconstruction of sparse signals from distorted randomized measurements," in *2010 IEEE International Conference on Acoustics Speech and Signal Processing (ICASSP)*. IEEE, 2010, pp. 3998–4001.
- [28] Z. Yang, L. Xie, and C. Zhang, "Off-grid direction of arrival estimation using sparse bayesian inference," *Arxiv preprint, available online at <http://arxiv.org/abs/1108.5838>*, 2011.
- [29] D. Wipf and B. Rao, "An empirical bayesian strategy for solving the simultaneous sparse approximation problem," *IEEE Transactions on Signal Processing*, vol. 55, no. 7, pp. 3704–3716, 2007.

- [30] M. Jordan, Z. Ghahramani, T. Jaakkola, and L. Saul, “An introduction to variational methods for graphical models,” *Machine learning*, vol. 37, no. 2, pp. 183–233, 1999.
- [31] D. Tzikas, A. Likas, and N. Galatsanos, “The variational approximation for bayesian inference,” *Signal Processing Magazine, IEEE*, vol. 25, no. 6, pp. 131–146, 2008.
- [32] B. Jørgensen, *Statistical properties of the generalized inverse Gaussian distribution*. Springer New York, 1982.
- [33] E. Candès and J. Romberg, “ ℓ_1 -magic: Recovery of sparse signals via convex programming,” *available online at <http://users.ece.gatech.edu/~justin/l1magic/downloads/l1magic.pdf>*, 2005.
- [34] M. Grant and S. Boyd, “CVX: Matlab software for disciplined convex programming,” *Available online at <http://cvxr.com/cvx>*, 2008.